

PATENT APPLICATION

**METHOD AND SYSTEM FOR DETERMINING THE POSITION OF A
BAD LINK**

Inventor(s): MICHAEL L. REGAL, a citizen of the United States of America, residing at
324 Brianne Court, Pleasanton, California 94566.

Assignee: CISCO TECHNOLOGY, INC. (A California Corporation)

Entity: LARGE

BACKGROUND OF THE INVENTION

5 [01] A ring of devices, such as port ASICs or stacked network devices, can be configured to operate together to aggregate functionality. For example, routers can be stacked so that they appear to a user as a single device having the aggregate functionality of the stacked individual devices.

 [02] A supervisory functionality must be defined to manage the individual
10 devices to transfer packets between ports on different devices and to perform other functions. The supervisory functionality can detect whether one of the links between devices or ports is not functional but the identity of the non-functional link must be determined by inspection.

 [03] For example, if the devices are port ASICs mounted on a PCB, then a physical hardware probe is used to determine the identity of the link that is not functioning
15 properly. The use of a hardware probe is time consuming and expensive since it requires disassembly of the box and a skilled technician to perform the test.

 [04] Accordingly, new techniques for identifying faulty links in a ring of devices are important to increase efficiency.

BRIEF SUMMARY OF THE INVENTION

20 [05] In one embodiment of the invention, at startup a device transmits a 'position' message to the downstream device connected to its transmitter link. A device's transmitted 'position' value is based on the 'position' value in the message received from the upstream device connected to its receiver link.

 [06] In another embodiment of the invention, if no 'position' message is
25 received from an upstream device, a device will advertise its 'position' as zero to its downstream link partner. If a 'position' message is received, a device will advertise its 'position' as being equal to the received message's 'position + 1'.

 [07] In another embodiment of the invention, a supervisor receives a 'position' message from the most downstream device and compares it to an expected position
30 number for a fully operational ring. If the supervisor determines that an unexpected value has been received, it can identify the physical location of the broken link as being equal to the expected value less the 'position' number received.

[08] In another embodiment of the invention, the physical position of a broken link is found by passing a 'position' message along a path of operating links. The 'position' value received by a supervisory function can compare the received 'position' value against an expected value to immediately determine the first physically broken link.

5 [09] Other features and advantages of the invention will be apparent in view of the following detailed description and appended drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[10] Fig. 1 is block diagram of a management ring;

10 [11] Fig. 2 is table depicting the format of an INIT command;

[12] Fig. 3 is flow chart depicting the behavior of the supervisory ASIC during INIT; and

[13] Fig. 4 is flow chart depicting the behavior of the port ASIC during INIT.

15

DETAILED DESCRIPTION OF THE INVENTION

[14] Reference will now be made in detail to various embodiments of the invention. Examples of these embodiments are illustrated in the accompanying drawings. While the invention will be described in conjunction with these embodiments, it will be understood that it is not intended to limit the invention to any embodiment. On the contrary, it is intended to cover alternatives, modifications, and equivalents as may be included within the spirit and scope of the invention as defined by the appended claims. In the following description, numerous specific details are set forth in order to provide a thorough understanding of the various embodiments. However, the present invention may be practiced without some or all of these specific details. In other instances, well known process operations have not been described in detail in order not to unnecessarily obscure the present invention.

[15] An embodiment of the invention will now be described, by way of example, not limitation, implemented in the system depicted in Fig. 1. Fig. 1 depicts a management ring 10 including a Supervisory ASIC (DEV0) 12 and first, second, and third port ASICs (DEV1, DEV2, and DEV3) 14a-c disposed on a platform.

30

[16] Each of the ASICs of Fig. 1 includes a management interface controller (MIC) including a receiver and a transmitter. In Fig. 1 the MIC receiver of DEV1 is coupled to the serial output of the supervisory ASIC (DEV0) by a first serial link 18a, the

MIC transmitter of DEV1 is coupled to the MIC receiver of DEV2 by a second serial link 18b, the MIC transmitter of DEV2 is coupled to the MIC receiver of DEV3 by a third serial link 18c, and the MIC transmitter of DEV3 is coupled to the serial input of the supervisory ASIC (DEV0) by a fourth serial link 18d.

5 [17] The MIC of a port ASIC includes a pass-through mechanism for rapidly forwarding messages from the MIC receiver to the MIC transmitter as soon as a determination is made that the port ASIC is not the target of the message.

 [18] In this embodiment, the bad link detection and identification process is performed during an initialization (INIT) process that can be triggered by a hard or soft reset,
10 by an error condition, or by other means.

 [19] The platform includes a non-volatile memory device, for example an EEPROM, holding information about the platform including the number of port ASICs present on the platform. During the INIT process, the supervisory ASIC reads the non-volatile memory device to determine the number of port ASICs on the platform that should
15 be functioning properly. In this example, counting begins at "0" so if the Supervisory ASIC reads "3" from the EEPROM it determines there are "4" ASICs on the board.

 [20] When the INIT process is triggered the supervisory ASIC outputs INIT. A special INIT message is used during the initialization process having a format depicted in Fig. 2. Referring to Fig. 2, the data format of the INIT includes a command field
20 that identifies the format as a management ring initialize command and a device number field which is used to address the port ASICs and to determine the presence and identity of a bad link.

 [21] The behavior of the supervisory ASIC and port ASICs during the INIT process is depicted in the flow charts of Fig. 3 and 4.

25 [22] Referring to Fig. 3, during INIT the supervisory ASIC starts a timer, outputs an INIT command having an device number field holding the value '0', and waits for an INIT command to be received. If an INIT command is not received before the timer times out a received new link position field is stored (see below) and an INIT failure is recorded.

 [23] Referring to Fig. 4, during INIT, until an INIT command is received a
30 Port ASIC continuously forwards OUT = 'LINK' to its MIC Transmitter, advertising a new Link Position of "0". If its MIC Receiver link is operational, the Port ASIC continuously receives IN = 'LINK' from an upstream Port ASIC, advertising its new Link Position of "0". In response to receiving these advertisements, the Port ASIC increments and continuously

forwards a modified new Link Position of 1 to its MIC Transmitter. This continues indefinitely, until the Port ASIC receives an 'INIT' on its MIC Receiver.

[24] As an example, in a system with 6 Port ASICs and no broken links, the Port ASIC sharing a link with the Supervisor ASIC's MIC Receiver link would advertise a new Link Position of "5"; the one Port ASIC sharing a link with the Supervisor ASIC's MIC Transmitter link would advertise a new Link Position of "0".

[25] When the INIT command is received at the MIC receiver of a port ASIC the port ASIC records the value held in the device number field of the received INIT command as its ID, increments the value held in the device number field, and outputs a modified INIT command having an device number field holding the incremented value.

[26] If each serial link is functioning properly, the supervisory ASIC outputs an INIT command having a device number field holding a value of '0'. This value is incremented to '1' by DEV1, to '2' by DEV2, and to '3' by DEV3. Thus, if every link on the ring is functioning then the supervisory receives an INIT command having an device number field having a value of '3' which corresponds to the number of devices that the supervisory ASIC knows are on the platform from having read the storage device on the platform.

[27] The case where the second serial link 18b is defective will now be described. DEV2 does not receive a valid INIT or LINK command. In this case, DEV2 will output a LINK command having a new link position field holding a value '0' since no valid command is received at its MIC receiver. This value will be incremented to '1' by DEV3 so that a LINK command having a new link position field holding the value '1' is received at the supervisor ASIC (DEV0).

[28] Because the second serial link 18b is defective, no INIT command will ever be received at the input of the supervisory ASIC so that the timer will time out and an INIT failure will be recorded and supervisory ASIC knows that there is a defective link in the management ring. Also, the value held in the stored new link field can be examined, in this case the value is "1", to determine that the second serial link 18b is the defective link.

[29] In this example, a new link field value of "1" indicates that the second serial link 18b is defective, and a new link field value of "0" indicates that third serial link 18c is defective. If all links were good the supervisory ASIC would receive "good value" of '3'. If one of the links is bad the location of the bad link is indicated by the difference between the good value and the received value. In the above example, if the second link 18b is bad, link value "1", then the received value of "2" identifies the second link 18b.

[30] If the fourth link 18d were defective the supervisory ASIC would receive an invalid command and therefore know that the fourth link 18d coupled to its input is defective.

5 [31] Other embodiments of the invention are not limited to ASICs on a single platform. For example, a stack of network devices, e.g., routing platforms, can be configured as a ring and one of the devices designated to function as a supervisor or start-up device. The processes described above would then be utilized to determine whether one of the links is defective and to identify which link is defective.

10 [32] Alternatively, the system could be configured to include network devices and/or port ASICs on the devices.

[33] As is known in the art, the mechanisms for implementing the various steps may be embodied in hardware or may be implemented as program code, stored on a computer readable medium, that is executed by a digital computer. The computer readable medium may include, among other things, magnetic media, optical media, electro-magnetic fields encoding digital information, and so on.

15 [34] The invention has now been described with reference to the preferred embodiments. Alternatives and substitutions will now be apparent to persons of skill in the art. Accordingly, it is not intended to limit the invention except as provided by the appended claims.

20